EXPLAINING THE RIDDLE OF DOLLAR COST AVERAGING

Abstract:

Dollar cost averaging (DCA) remains a very popular investment strategy, even though previous academic research has long shown that in normal circumstances it is mean-variance inefficient. More recent research explains DCA's continued popularity by assuming that investors are subject to behavioral finance effects or have non-variance risk preferences. The present paper shows that DCA's popularity can instead be explained by a specific and demonstrable cognitive error. This is a simpler hypothesis since it requires no additional assumptions about investor preferences. Identifying this error should also help investors make much better informed decisions about whether to use DCA.

This version: 14 June 2010
Comments welcome

Simon Hayley
Lecturer in Finance
Cass Business School
simon.hayley.1@city.ac.uk

Acknowledgements: I am grateful to George Constantinides, Christopher Gilbert, Stewart Hodges, Burton Malkiel, Hugh Patience and Meir Statman for comments. The usual disclaimer applies.

EXPLAINING THE RIDDLE OF DOLLAR COST AVERAGING

Dollar cost averaging (DCA) is the practice of building up investments gradually over time in equal dollar amounts, rather than investing the desired total in one lump sum. Proponents of DCA argue that as it reduces the average cost of investing (since more securities are purchased in periods when the price is relatively low), it must generate higher returns.

By contrast, previous academic research has established that in normal circumstances DCA is not mean-variance efficient. Despite this, DCA remains very popular among investors, and it is still recommended by influential authors and commentators. Some funds explicitly encourage investors to make regular contributions as part of a DCA strategy. Many investors have regular savings plans, and hence are in effect using a DCA strategy, but proponents of DCA go beyond this and argue that cash already saved should be invested gradually. This paper investigates why the views of proponents of DCA are so much at odds with the academic findings.

Constantinides [1979] showed that as DCA allows no flexibility once the investor has started the planned series of investments, so this strategy must be dominated by others which allow the investor to make use of the additional information which is available in later periods. Empirical studies also come out in favor of investing the whole desired amount in one lump sum rather than waiting. The strategy which turns out to have been better over any given period depends on the particular path taken by prices, but a large number of studies has found that investing in one lump sum has generally given better mean-variance performance than DCA. These include Knight and Mandell [1992/93], Williams and Bacon [1993], Rozeff [1994] and Thorley [1994]. DCA may seem to improve diversification by making many small purchases but, as Rozeff notes, the result is that overall profits are most sensitive to returns in the later part of the period, when the investor is nearly fully invested. Earlier returns are given correspondingly little weight, since the investor then holds mainly cash. Better diversification is achieved by investing in one initial lump sum, and thus being equally exposed to the returns in each sub-period.

Milevsky and Posner [1999] extend the analysis into a continuous time framework, and show that it is always possible to construct a constant proportions continuously rebalanced portfolio which will stochastically dominate DCA in a mean-variance framework. They also show that for typical levels of volatility and drift there will be a static buy and hold strategy which dominates DCA.

As the evidence became overwhelming that DCA does not have the benefits claimed for it, research turned increasingly to attempts to explain why investors nevertheless persist with it. Statman [1995] showed that behavioral finance can explain DCA’s continued popularity: the strategy helps investors frame decisions in an artificially flattering context. It also commits them to continue investing at a constant rate, allowing limited choice in the short term. This reduces regret and protects investors from their tendency to base investment decisions on naïve extrapolation of recent price trends.

Milevsky and Posner [1999] also show that the expected payoff from DCA can be attractive if investors have a fixed target portfolio value. More recent studies have shown that DCA can generate attractive reductions in some non-variance measures of risk (Dubil [2005], Trainor [2005]), but not others (Leggio and Lien [2003]).

The academic research in this field has become steadily more sophisticated, yet DCA is generally still recommended to investors on the basis of a simple demonstration that DCA builds up investments at a cost which is below the average price. The continued popularity of DCA suggests that investors remain so convinced that buying at a lower average cost must increase profits that they are impervious to academic arguments to the contrary. Previous authors have noted that the ability to purchase shares at below their average price is irrelevant, since investors cannot subsequently sell at this average price (e.g. Thorley [1994] and Milevsky and Posner [1999]), but this does not address the argument directly, since proponents are not arguing that DCA is a means of making guaranteed short-term profits, but merely that it is a better way to enter longer-term trades.

Unlike previous research, the present paper focuses directly on the argument put forward by proponents of DCA. It finds that it is extremely misleading to compare the average cost achieved by DCA with the average price: this implicitly compares DCA with a strategy which uses perfect foresight to invest more when prices are about to fall and less when they are about to rise. It is only because of this that DCA appears to offer higher returns.

This result achieves two things. First, a positive point: the continued popularity of DCA can now be explained as resulting from a specific and demonstrable cognitive error. This is a simpler hypothesis than those offered by previous studies, which can only explain DCA's popularity by assuming the existence of additional investor characteristics which cannot be directly observed. Second, a normative point: exposing the flaw in the argument put forward by proponents of DCA, will help investors make more informed decisions about whether to use the strategy.

The structure of this paper is as follows. The next section analyses the implicit counterfactuals that are being compared when we note that DCA purchases shares at below the average price, and gives an intuitive demonstration of exactly why this does not imply that DCA will generate higher profits. The section after that demonstrates this result more formally. Conclusions are drawn in the final section.

The Intuition

Table 1 shows a numerical example typical of those used by proponents of DCA. The first columns show the purchases made under a DCA strategy (the alternative ESA strategies are not normally made explicit and will be explained later). A fixed $60 each period is invested in a specific equity. The share price is initially $3, allowing 20 units to be purchased. The sharp fall in price to $1 allows 60 units to be purchased for the same dollar outlay in period 2, whilst the rebound to $2 allows 30 units to be bought in the final period.

The argument in favor of DCA is that it buys shares at an average cost ($180/110 = $1.64) which is lower than the average market price of the shares over the period during which they were accumulated ($2). This is achieved because DCA automatically buys more shares during periods when they are relatively cheap and fewer when they are more expensive.

Table 1		(a) DCA		(b) ESA1		(c) ESA2	
Period	Share price	Shares purchased	Investment	Shares purchased	Investment	Shares purchased	Investment
1	$3	20	$60	20	$60	30	$90
2	$1	60	$60	20	$20	30	$30
3	$2	30	$60	20	$40	30	$60
Total		110	$180	60	$120	90	$180

Greenhut [2006] takes issue with the particular return assumptions used in some such "demonstrations" of the superiority of DCA. However, there is a much more general issue here: the average purchase cost for DCA investors gives greater weight to periods when the price is relatively low, so price fluctuations will always mean that DCA investors buy at less than the average price, regardless of the particular path taken by prices. The difference is particularly large in the example above due to the extreme price volatility, but any price volatility favors DCA: only when the share price remains unchanged in all periods will the average cost equal the average price. Rather than challenging the particular numbers used, we need to examine why it is that a strategy which appears to buy assets at a lower average cost does not in fact lead to higher expected profits.

As described above, previous studies have found that DCA is mean-variance inefficient compared to investing the whole desired amount immediately in one lump sum. But proponents of DCA are making a different comparison. In noting that the average cost achieved under DCA is less than the average price they are implicitly comparing DCA with a strategy of buying a fixed number of shares in each period. This is the comparison that we will make here in order to understand why the case in favor of DCA is misleading.

Table 1 compares the cashflows under DCA with two alternative strategies which buy a constant number of shares in each period (equal share amounts: ESA1 and ESA2). The difference between these two alternatives may appear to be a trivial matter of scale, but it is in this difference that the fallacy underlying DCA lies.

ESA1 is an attempt to invest the same total amount as DCA over these three periods, but to do so in equal share amounts. With the share price initially at $3, a reasonable approach would be to buy 20 shares, since if prices remain at this level in periods 2 and 3 we will end up investing exactly the $180 total that we desire. But our strategy then requires that we buy 20 shares in each of the following periods, and when prices in periods 2 and 3 turn out to be substantially lower, we end up investing only $120: much less than we had intended. It is only with perfect foreknowledge of future share prices that we could have known that the only way of investing $180 in equal share amounts is to buy 30 shares each period, as shown in ESA2.

The fact that DCA buys shares at an average cost which is below the unweighted average price during this period effectively compares the DCA strategy with the ESA2 strategy which invests the same dollar total in equal share amounts. But, as we have seen, ESA2 can only achieve this if we have perfect foresight – otherwise we will generally end up investing the wrong amount.

Moreover, this foresight is used in a way which systematically reduces profitability. In this example, the ESA2 strategy reacts to the knowledge that prices are about to fall by investing more than it otherwise would in period 1. Conversely, it would invest less in period 1 if prices in subsequent periods were going to be higher. This is the only way to invest the correct amount, but it is, of course, systematically loss-making behavior.

Table 2 shows the same strategies, but with the share price rising rather than falling. The DCA strategy again invests $60 each period, but as prices rise fewer shares are purchased in the later periods. Once again DCA achieves an average cost ($180/47=$3.83) below the average price ($4) by buying more shares when they are relatively cheap. Correspondingly, the ESA2 strategy invests the same total amount, but buys only 45 shares compared to 47 using DCA.

Table 2		(a) DCA		(b) ESA1		(c) ESA2	
Period	Share price	Shares purchased	Investment	Shares purchased	Investment	Shares purchased	Investment
1	$3	20	$60	20	$60	15	$45
2	$4	15	$60	20	$80	15	$60
3	$5	12	$60	20	$100	15	$75
Total		47	$180	60	$240	45	$180

However, as we saw earlier, the real alternative to DCA is ESA1. In practice our best guess would again be to invest one third of our total budget in the first period, since if prices were to stay at this level we would invest the correct amount. But when prices subsequently rise we end up spending substantially more than this ($240). ESA2 invests the correct amount, but it achieves this only by knowing that prices are about to rise, and responding to this knowledge by buying fewer shares than ESA1. Again, profits are reduced.

This shows us the flaw in the argument put forward by proponents of DCA. It is true that DCA allows us to purchase shares at below their average price, but this in effect compares DCA with a strategy which uses perfect foresight to systematically reduce profits and increase losses. The following section demonstrates this result formally and shows that it is only for this reason that DCA appears to offer superior profits.

The Proof

We consider investing over a series of *n* discrete periods. The price of the asset in each period *i* is p_i. The alternative investment strategies differ in the quantity of shares q_i that are purchased in each period. We evaluate profits at a subsequent point, after all investments have been made. If prices are then p_T, the profit made by any investment strategy is:

$$\Pi = p_T \sum_{i=1}^{n} q_i - \sum_{i=1}^{n} p_i q_i \tag{1}$$

We define DCA as a strategy which invests *b* dollars in each period $(p_i q_i = b \quad \forall i)$. This gives us the profits that will result from following a DCA strategy:

$$\Pi_{dca} = p_T \sum_{i=1}^{n} \left(\frac{b}{p_i} \right) - nb \tag{2}$$

We assume that investors do not believe that they can forecast market prices - in effect they assume that prices follow a random walk. However, we should stress that this is a statement about investors' ex ante expectations, and does not imply any presumption that markets are in fact weak form efficient. The key point in this context is that DCA will only ever be an attractive strategy for investors who do not believe that they can forecast short-term price movements. DCA commits investors to invest the same amount no matter what price movements they expect in the coming period – those who feel that they can forecast short-term price movements will reject this and follow other strategies instead.

We also assume that this random walk has zero drift. This assumption is also a statement about investors' ex ante expectations rather than about markets themselves. Investors presumably believe that over the medium term their chosen securities will generate an attractive return, but they must also believe that the return over the short term (while they are using DCA to build up their position) is likely to be small. Investors who expect significant returns over the short term would prefer to invest immediately in one lump sum rather than delay their investments by following a DCA strategy.

The assumption of zero drift need not imply a loss of generality, since drift could be incorporated into this framework by defining prices not as absolute market prices, but as prices relative to a numeraire which appreciates at a rate which gives a fair return for the risks inherent in this asset ($p_i^* = p_i/(1+r)^i$, where *r* reflects the cost of capital and a risk premium appropriate to this asset). We could then assume that p_i^* has zero expected drift since investors who use DCA will not believe that they can forecast short-term relative returns for assets of equal risk: those who do would again reject trading strategies which forced them to delay their purchases. The results derived below would continue to hold for p_i*, with profits then defined as excess returns compared to the risk-adjusted cost of capital.[1]

Given these assumptions, investors will assume ex ante that prices will remain flat, with $E[p_T/p_i]=1$ for all *i*. Substituting this into equation 2, we see that the ex ante expected profit from our DCA strategy is zero.

Our alternative investment strategy is to buy equal numbers of shares in each period $(q_i = a \quad \forall i)$. Substituting this into equation 1 gives us:

$$\Pi_{esa1} = anp_T - a\sum_{i=1}^{n} p_i \qquad (3)$$

This shows that the ex ante expected profit from our ESA strategy is also zero (this can be seen by substituting $E[p_i]=E[p_T]$ for all *i*, as an equivalent expression of our driftless random walk). Thus DCA does not give superior expected returns.

This is an intuitive result. We can regard our total return as a weighted average of the returns made on the amounts invested in each period. ESA and DCA differ only in giving different relative weights to these individual period returns. But if prices are believed to follow a random walk with zero drift the expected return will be zero for each period and varying the relative weight on different periods' returns cannot change the expected aggregate return. By contrast, DCA's supporters suggest that even when investors have no belief that they can forecast market returns they can nevertheless expect to beat the market when using DCA.

As we saw in the previous section, the total amount invested under ESA1 ($a\sum p_i$) is likely to differ from the amount (*nb*) invested under DCA. But the comparison that is usually presented by proponents of DCA assumes that the two techniques invest equal total amounts. Thus in order to duplicate the conventional "proof" of the benefits of DCA, we need to rescale the number of shares bought under ESA1 by the fixed factor ($nb/a\sum p_i$), so that an exactly equal amount is invested by the two strategies. This gives us the expected profits resulting from strategy ESA2:

$$\Pi_{esa2} = \Pi_{esa1}\left(\frac{nb}{a\sum_{i=1}^{n} p_i}\right) \qquad (4)$$

The use of foresight can be seen in the fact that the scaling factor depends on the average share price throughout the investment period. Only if this is known at the outset would we be able to buy the correct number of shares so that we end up spending exactly the same amounts under ESA2 and DCA. Substituting from equation 3:

$$\Pi_{esa2} = \left(anp_T - a\sum_{i=1}^{n} p_i\right)\left(\frac{nb}{a\sum_{i=1}^{n} p_i}\right) \qquad (5)$$

[1] We must also assume that funds not yet needed can be held in assets with the same expected return. This assumption is clearly generous to DCA – if instead cash is held on deposit at lower expected return, then DCA's expected return will clearly be reduced by delaying investment.

$$= \frac{bn^2 p_T}{\sum_{i=1}^{n} p_i} - nb \qquad (6)$$

Subtracting equation (6) from equation (2) we find:

$$\Pi_{dca} - \Pi_{esa2} = nbp_T \left(\frac{1}{n} \sum_{i=1}^{n} \left(\frac{1}{p_i} \right) - \frac{n}{\sum_{i=1}^{n} p_i} \right) \qquad (7)$$

The term in brackets is non-negative for positive p_i, and strictly positive if they are not all equal. This follows from directly from the arithmetic mean–harmonic mean inequality[2].

This achieves our objective. We have shown that the expected profits from a DCA strategy are identical to those from our ESA1 strategy (both give zero expected profits). By contrast, DCA gives higher expected profits than our ESA2 strategy which scales the level of investment so as to spend exactly the same total amount as DCA. However, ESA2 is not a feasible strategy, since it uses perfect foresight to invest in a systematically loss-making fashion. It is only on this biased comparison that DCA appears to make greater returns, yet it is exactly this comparison which is implicitly being made when we note that DCA buys at lower average cost than the average price.

---

1. The arithmetic-harmonic mean inequality is usually stated as: $(x_1 + ... + x_n)/n \geq \frac{n}{\frac{1}{x_1} + ... + \frac{1}{x_n}}$

for positive x_i, so we have substituted $x_i = \frac{1}{p_i}$. This inequality follows directly from Jensen's inequality that $E[f(x)] \geq f(E[x])$ for any convex function $f(.)$, using the function $f(x) = 1/x$.

Conclusion

DCA remains a very popular strategy even though previous research has demonstrated that under normal circumstances it is mean-variance inefficient. More recent research has shown that behavioral factors and non-variance risk preferences can be used to explain why investors still use this strategy. However, by far the most common argument put forward in favor of DCA is that it must generate higher profits since it allows purchases to be built up over time at an average cost which is below the average market price. The central role that is given to this argument suggests that it is a key factor underpinning DCA's continued popularity.

This paper tackles this argument head on, and finds that the higher profits that are claimed for DCA result from an implicit comparison with a counterfactual strategy which uses perfect foresight in a way that consistently reduces profits, systematically investing more when prices are about to fall, and less when they are about to rise. DCA offers no additional profits when compared with a sensible alternative strategy. This biased counterfactual is buried in what otherwise seems the obvious and uncontentious comparison of the average purchase cost achieved by DCA with the average market price over the period.

Identifying this hidden bias means that we can now regard the continued popularity of DCA as resulting from a specific and demonstrable cognitive error. This explanation uses the observable fact that proponents of DCA almost invariably use a flawed argument to suggest that DCA generates higher profits. By contrast, behavioral finance effects can only explain DCA's popularity by assuming the existence of additional investor characteristics which cannot be directly observed. Behavioral finance effects remain very plausible, but Occam's razor would suggest that the simpler hypothesis should be preferred.

There is also a normative aspect. Exposing the flaw in the argument put forward by proponents of DCA should help investors make more informed decisions about whether to use the strategy. If, aside from this cognitive error, investors are rational agents who optimize in a mean-variance framework, then they will abandon DCA and switch to strategies which are more mean-variance efficient. This would be unambiguously welfare-improving.

Previous research has shown that if investors have non-variance risk preferences, they may prefer to continue using DCA. Behavioral finance effects such as the avoidance of regret can also bring real psychological benefits. However, abandoning a misguided belief that DCA increases expected returns will leave investors better able to assess whether such non-pecuniary benefits and alternative risk preferences justify the use of DCA. The existence of such effects does not alter our conclusion that eliminating the cognitive error should improve investor welfare.

It could be argued that DCA would also bring genuine financial benefits if it helps to save investors from a tendency to naively extrapolate market trends (reducing their bias towards buying at the top of upswings and selling at the bottom of downturns). This argument would suggest that if investors are currently subject to two offsetting cognitive errors, eliminating just one of them (the belief that DCA raises expected returns) could reduce welfare by leaving them more at risk of making other damaging mistakes. However, this line of reasoning would leave academics in the untenably paternalistic position of preferring to leave investors believing a demonstrable falsehood. We should instead seek to publicize both errors. Improved investor education is a valid goal, and selectively withholding information is not a defensible means of improving investor welfare.

References

Constantinides, G.M. “A note On The Suboptimality Of Dollar-Cost Averaging As An Investment Policy.” Journal of Financial and Quantitative Analysis, Vol.14, No. 2 (June 1979), pp. 443-450.

Dubil, R. “Lifetime Dollar-Cost Averaging: Forget Cost Savings, Think Risk Reduction.” Journal of Financial Planning, Vol. 18, Issue 10 (October 2005), pp. 86-90.

Greenhut, J.G. “Mathematical Illusion: Why Dollar-Cost Averaging Does Not Work.” Journal of Financial Planning, Vol. 19, Issue 10 (October 2006), pp. 76-83.

Knight, J.R. and Mandell, L. “Nobody Gains From Dollar Cost Averaging: Analytical, Numerical And Empirical Results.” Financial Services Review, Vol. 2, Issue 1 (1992/93), pp. 51-61.

Leggio, K. and Lien, D. “Comparing Alternative Investment Strategies Using Risk-Adjusted Performance Measures.” Journal of Financial Planning, Vol. 16, No. 1 (January 2003), pp. 82-86.

Milevsky, M.A. and Posner, S.E. "A Continuous-Time Re-examination of the Inefficiency of Dollar-Cost Averaging." International Journal of Theoretical & Applied Finance, Vol. 6, Issue 2 (March 2003), pp. 173-194.

Rozeff, M.S. “Lump-sum Investing Versus Dollar-Averaging.” Journal of Portfolio Management, Vol. 20, Issue 2 (winter 1994), pp. 45-50.

Statman, M. “A Behavioral Framework For Dollar-Cost Averaging.” Journal of Portfolio Management, Vol. 22, No. 1 (Fall 1995), pp. 70-78.

Thorley, S.R. “The Fallacy of Dollar Cost Averaging.” Financial Practice and Education, Vol. 4, No. 2 (Fall/Winter 1994), pp. 138-143.

Trainor, William J Jr. “Within-horizon exposure to loss for dollar cost averaging and lump sum investing.” Financial Services Review; Vol. 14, Issue 4 (Winter 2005), pp. 319-330.

Williams, R.E. and Bacon, P.W. “Lump-sum Beats Dollar Cost Averaging.” Journal of Financial Planning, Vol. 6, No.2 (April 1993), pp. 64–67.